arXiv:1602.08521v1 [stat.ME] 26 Feb 2016

Theoretical Properties and Practical Performance of Fully Robust One-Sided Cross-Validation

Olga Y. Savchuk, Jeffrey D. Hart

Abstract

Fully robust OSCV is a modification of the OSCV method that produces consistent bandwidth in the cases of smooth and nonsmooth regression functions. The current implementation of the method uses the kernel H_I that is almost indistinguishable from the Gaussian kernel ϕ on the interval $[-4, 4]$, but has negative tails. The theoretical properties and practical performances of the H_I- and ϕ-based OSCV versions are compared. The kernel H_I tends to produce too low bandwidths in the smooth case. The H_I-based OSCV curves are shown to have wiggles appearing in the neighborhood of zero. The kernel H_I uncovers sensitivity of the OSCV method to a tiny modification of the kernel used for the cross-validation purposes. The recently found robust bimodal kernels tend to produce OSCV curves with multiple local minima. The problem of finding a robust unimodal nonnegative kernel remains open.

Keywords: cross-validation; one-sided cross-validation; local linear estimator; bandwidth selection; mean average squared error.

AMS Subject Classifications: 62G08; 62G20.

1 Introduction

Nonparametric regression estimation involves selecting a smoothing parameter, usually called the *bandwidth*, that mainly determines the appearance of a regression estimate. Inappropriately small bandwidth results in a bumpy estimate that tracks almost every data point on the scatter diagram, whereas too large bandwidth produces an oversmoothed regression estimate that may fail to represent important features of the regression function such as multiple peaks, sharpness of a peak, etc. There exist many methods that use the data to estimate the bandwidth that is optimal in certain sense. The most frequently used data-based bandwidth selection methods are the plug-in rule of Ruppert et al. (1995) and the cross-validation (CV) method of Stone (1977). There are many variations of both methods. One of the successful modifications of the CV method is the one-sided cross-validation (OSCV) method developed by Hart and Yi (1998).

All original OSCV research relies on the assumption that the regression function is *smooth*, which means that it has at least two continuous derivatives. Hart and Yi (1998) showed that using OSCV instead of CV may produce up to twentyfold reductions of the asymptotic bandwidth variance. Yi (2005) conducted a simulation study to illustrate the improved stability of OSCV compared to CV in finite samples. Hart and Lee (2005) argued

that OSCV is robust to moderate levels of autocorrelation. Martínez-Miranda et al. (2009) developed a version of the OSCV method for the kernel density estimator.

For many real data sets in economics, medicine, biology and other fields, the relationship between the variables is described by a continuous function that has sharp corners or *cusps* appearing at the points where the first derivative of a function has simple discontinuities. A continuous regression function with cusps is refereed to as *nonsmooth*. The original OSCV method produces a biased estimator of the optimal bandwidth in the nonsmooth case. Savchuk et al. (2013) developed the *fully robust OSCV* version that results in a consistent estimation of the optimal bandwidth regardless of the regression function's smoothness.

This article provides a detailed investigation of the theoretical properties and practical performance of the current implementation of the fully robust OSCV method. We also demonstrates performance of OSCV based on our recently found robust bimodal kernel.

The rest of the article proceeds as follows. In Section 2 we overview the problem of nonparametric regression estimation and outline the steps in the OSCV method. Section 3 contains an extended discussion of the fully robust OSCV method and brings new light onto performance of the original OSCV version of Hart and Yi (1998). Section 4 contains summary of our findings. The appendix includes certain supplementary materials.

Our subsequent presentation requires introducing the following notation. For an arbitrary function g, define

$$R_g = \int_{-\infty}^{\infty} g^2(u)\,du, \qquad \mu_{2g} = \int_{-\infty}^{\infty} u^2 g(u)\,du, \tag{1}$$

$$J_g = \left(\frac{R_g}{\mu_{2g}^2}\right)^{1/5}, \tag{2}$$

$$B_g = \int_0^1 \left\{z\big(1 - D_g(z)\big) + G_g(z)\right\}^2 dz + \int_0^1 \left\{zD_g(-z) + G_g(-z)\right\}^2 dz, \tag{3}$$

and for all z,

$$D_g(z) = \int_{-\infty}^{z} g(u)\,du,$$
$$G_g(z) = \int_{-\infty}^{z} ug(u)\,du.$$

2 Nonparametric regression estimation and the OSCV method

In the nonparametric regression model the observations $Y_1, Y_2, \ldots, Y_n$ are assumed to be generated as

$$Y_i = r(x_i) + \varepsilon_i,$$

where r is an unknown regression function defined on the interval $[0, a]$, $a > 0$, and $\varepsilon_1, \varepsilon_2, \ldots, \varepsilon_n$ are uncorrelated error terms such that $E(\varepsilon_i) = 0$ and $\mathrm{Var}(\varepsilon_i) = \sigma^2$, $i = 1, \ldots, n$. The design points $x_1 < x_2 < \cdots < x_n$ are assumed to be fixed quantiles of the design density f. In the case of an evenly spaced design, $f \equiv 1/a$.

The OSCV method is intended to select the bandwidths of the Gasser-Müller estimator (see Gasser and Müller (1979)) or the local linear estimator (see Cleveland (1979)). In this

article we concentrate on the local linear estimator, defined at the point x as

$$\hat{r}_h(x) = \frac{\sum_{i=1}^n w_i(x) Y_i}{\sum_{i=1}^n w_i(x)}, \tag{4}$$

where $h > 0$ is the bandwidth,

$$w_i(x) = K\left(\frac{x - x_i}{h}\right)\left(t_{n,2} - (x - x_i)t_{n,1}\right), \tag{5}$$

and

$$t_{n,j} = \sum_{i=1}^n K\left(\frac{x - x_i}{h}\right)(x - x_i)^j, \quad j = 1, 2. \tag{6}$$

The kernel function K is of the second order, that is it integrates to one, has zero first moment, and finite second moment (see Wand and Jones (1995)).

Some popular measures of closeness of $\hat{r}_h$ to r are the mean average squared error (MASE) and the average squared error (ASE). The ASE function is defined as

$$\mathrm{ASE}_K(h) = \frac{1}{n}\sum_{i=1}^n \left(\hat{r}_h(x_i) - r(x_i)\right)^2.$$

The subscript "K" is used above to emphasize dependance of $\hat{r}_h$ on the kernel K. The bandwidth $\hat{h}_0$ that minimizes ASE is optimal for the data set at hand. The MASE function is defined as the expectation of the ASE function. The bandwidth h_0 that minimizes MASE is optimal in the average sense for all data sets generated from r at the fixed values of n and σ.

In the case of a smooth regression function, the asymptotic MASE expansion for the local linear estimator $\hat{r}_h$ based on the kernel K has the following form:

$$\mathrm{MASE}_K(h) = \mathrm{AMASE}_K(h) + o\left(h^4 + \frac{1}{nh}\right),$$

where

$$\mathrm{AMASE}_K(h) = \frac{R_K \sigma^2}{nh} + \frac{\mu_{2K}^2 h^4}{4}\int_0^a (r''(x))^2 f(x)\, dx. \tag{7}$$

The minimizer of the $\mathrm{AMASE}_K(h)$ function is

$$h_n = \left(\frac{R_K \sigma^2}{\mu_{2K}^2 \int_0^a \left(r''(x)\right)^2 f(x)\, dx}\right)^{1/5} n^{-1/5} = J_K C_{r,\sigma} n^{-1/5},$$

where

$$C_{r,\sigma} = \left(\frac{\sigma^2}{\int_0^a \left(r''(x)\right)^2 f(x)\, dx}\right)^{1/5}. \tag{8}$$

The OSCV method is designed to produce an estimate of the MASE-optimal bandwidth h_0. The main idea behind OSCV is to use different kernels in the estimation and cross-validation stages. The final regression estimate $\hat{r}_h$ is computed by using the local linear estimator based on a highly efficient kernel K, such as Gaussian, Epanechnikov, etc. Kernel

efficiency discussion may be found in Wand and Jones (1995). In the cross-validation stage, one uses a so-called one-sided estimator $\tilde{r}_b$, where $\tilde{r}_b(x_i)$ is a local linear estimator computed from the data points $(x_1, Y_1), \ldots, (x_i, Y_i)$, $i = 1, \ldots, n$. The estimator $\tilde{r}_b$ depends on the bandwidth b that is generally different from the bandwidth h used in $\hat{r}_h$. Moreover, $\tilde{r}_b$ is based on the kernel H that may differ from the kernel K used in $\hat{r}_h$. Thus, $K = H$ in the original OSCV implementation of Hart and Yi (1998), but $K \neq H$ in the fully robust OSCV method of Savchuk et al. (2013). In the asymptotic sense, computing $\tilde{r}_b$ by using the data on only one side of an estimation point is equivalent to using all data points in the local linear estimator based on the so-called one-sided kernel L related to the kernel H in the following way:

$$L(u) = 2H(u)\frac{S_2 - uS_1}{S_2 - 2S_1^2} I_{[0,\infty)}(u), \tag{9}$$

where

$$S_i = \int_0^\infty u^i H(u)\, du, \quad i = 1, 2,$$

and I_A is the indicator function of set A.

The one-sided estimator $\tilde{r}_b$ is used to compute the one-sided cross-validation function defined as

$$\mathrm{OSCV}(b) = \frac{1}{n-m} \sum_{i=m+1}^{n} (\tilde{r}_b^i(x_i) - Y_i)^2, \tag{10}$$

where $\tilde{r}_b^i$ is the leave-one-out version of $\tilde{r}_b$. Thus, $\tilde{r}_b^i(x_i)$ is computed from the observations $(x_1, Y_1), \ldots, (x_{i-1}, Y_{i-1})$. The quantity m is the number of the data points that are used to compute $\tilde{r}_b^{m+1}(x_{m+1})$. It is common to take $m = 4$. Let $\hat{b}_{OSCV}$ denote the minimizer of the OSCV function.

The OSCV function (10) is defined by analogy with the cross-validation function of Stone (1977) that is given by

$$\mathrm{CV}(h) = \frac{1}{n} \sum_{i=1}^{n} \left(\hat{r}_h^{-i}(x_i) - Y_i\right)^2.$$

In the above expression $\hat{r}_h^{-i}$ is the leave-one-out estimator that is the local linear estimator computed from all data except for the ith observation. Let $\hat{h}_{CV}$ denote the minimizer of the CV function.

Let AMASE_L denote the AMASE function for the local linear estimator based on the kernel L in the case when r is smooth. An expression for AMASE_L is obtained from (7) by everywhere replacing K by L. The minimizer of AMASE_L is denoted by b_n. It appears that

$$\frac{h_n}{b_n} = \left(\frac{R_K}{\mu_{2K}^2} \cdot \frac{\mu_{2L}^2}{R_L}\right)^{1/5} = \frac{J_K}{J_L} \equiv C, \tag{11}$$

where J_K and J_L are computed for the kernels K and L, respectively, according to (2). The constant C is referred to as the *smooth constant* and is completely determined by the kernels K and H.

Hart and Yi (1998) argued that in the case when r is smooth, the OSCV function is approximately unbiased estimator of $\mathrm{MASE}_L + \sigma^2$. This implies that $\hat{b}_{OSCV}$ estimates b_n. These considerations justify the following OSCV method's bandwidth selection rule:

$$\hat{h}_{OSCV} = C\hat{b}_{OSCV}, \tag{12}$$

where C is defined by (11). It appears that $\hat{h}_{OSCV}$ is a consistent estimator of h_0 in the case when r is smooth. The OSCV regression estimate is computed as the local linear estimate (4) based on the bandwidth $\hat{h}_{OSCV}$.

Savchuk et al. (2013) developed the OSCV theory in the case when the regression function r is nonsmooth. Given that the derivative of r has jumps at the points $\{x^{(t)}\}$, $t = 1, \ldots, k$, the asymptotic MASE expansion for the local linear estimator $\hat{r}_h$ based on the kernel K has the following form:

$$\mathrm{MASE}_K(h) = \mathrm{AMASE}^*_K(h) + O\left(h^4 + \frac{1}{n^2h^3}\right) + o\left(\frac{1}{nh}\right),$$

where

$$\mathrm{AMASE}^*_K(h) = \frac{R_K\sigma^2}{nh} + h^3 B_K \sum_{t=1}^{k} f(x^{(t)}) \left(r'(x^{(t)}+) - r'(x^{(t)}-)\right)^2. \tag{13}$$

The value of B_K is computed for the kernel K according to (3). Savchuk et al. (2013) derived the result (13) in the case of a regression function defined on $[0, 1]$, but it also holds in the case of a regression function defined on $[0, a]$. The minimizer of AMASE^*_K has the following form:

$$h^*_n = \left(\frac{R_K\sigma^2}{3B_K \sum_{t=1}^{k} f(x^{(t)})(r'(x^{(t)}+) - r'(x^{(t)}-))^2}\right)^{1/4} n^{-1/4}.$$

Let AMASE^*_L denote the AMASE function that is computed for the local linear estimator based on the kernel L in the case when r is nonsmooth. An expression for AMASE^*_L follows from (13) by everywhere replacing K by L. Let b^*_n denote the minimizer of AMASE^*_L. It follows that

$$\frac{h^*_n}{b^*_n} = \left(\frac{R_K}{B_K} \cdot \frac{B_L}{R_L}\right)^{1/4} \equiv C^*, \tag{14}$$

where R_K, B_K and R_L, B_L are computed for the kernels K and L, respectively, according to(1) and (3). The constant C^* is completely determined by the kernels K and H and is referred to as the *nonsmooth constant*.

In the nonsmooth case, the relative bandwidth bias increase due to inappropriate using C instead of C^* can be assessed as

$$E_C = \frac{C\hat{b}_{OSCV} - C^*\hat{b}_{OSCV}}{C^*\hat{b}_{OSCV}} \cdot 100\% = \frac{C - C^*}{C^*} \cdot 100\%. \tag{15}$$

It follows from the results of Savchuk et al. (2013), that $E_C < 7.01\%$ for such frequently used kernels as Epanechnikov and quartic. However, in the case of the Gaussian kernel ϕ, defined as $\phi(u) = (2\pi)^{-1/2}\exp(-u^2/2)$, the discrepancy $E_C = 16.74\%$. Indeed in the Gaussian case the smooth constant $C_\phi = 0.6168$, and the nonsmooth constant $C^*_\phi = 0.5284$. Simulation study of Savchuk et al. (2013) confirms that OSCV tends to produce too large bandwidths in the case when r is nonsmooth and $K = H = \phi$. Our experience with smoothing suggests that the relative bias of 7% has a negligible effect on performance of an estimator. However, the bias increase of 16% indicates that the methods requires a bias correction.

Theoretically, the bandwidth bias in the case when r is nonsmooth can be eliminated by replacing C by C^* in the OSCV bandwidth rule (12). However, such replacement should be justified by either prior information about nonsmoothness of r or the existence of cusps

in r should be evident from the scatter diagram of the data. The cusps may be masked by the noise in the data, so the analyst may erroneously apply a smooth version of the OSCV rule (12) to a nonsmooth function r.

Interestingly, the bandwidth bias introduced by inappropriate use of the smooth constant C in the nonsmooth case has a trivial effect on MASE, at least asymptotically. This is shown in Savchuk et al. (2013) based on the following measure of error:

$$E_{MASE} = \left(\frac{\mathrm{AMASE}_K^*(Cb_n^*)}{\mathrm{AMASE}_K^*(h_n^*)} - 1 \right) \cdot 100\%.$$

The bandwidth h_n^* is an asymptotic analog of the bandwidth $C^*\hat{b}_{OSCV}$, whereas the quantity Cb_n^* is asymptotically equivalent to the OSCV bandwidth (12) that is not justified for the case when r is nonsmooth. It appears that

$$E_{MASE} = \left(\frac{3}{4}x + \frac{1}{4x^3} - 1 \right) \cdot 100\%,$$

where

$$x = \left(\frac{B_L}{B_K} \right)^{1/4} \left(\frac{R_K}{R_L} \right)^{1/20} \left(\frac{\mu_{2K}^2}{\mu_{2L}^2} \right)^{1/5}.$$

The quantity E_{MASE} is 0.72% for the Epanechnikov kernel, it is 0.73% for the quartic kernel, but it is equal to 4.02% for the Gaussian kernel. Attempts to improve statistical properties of the OSCV method in the case when r is nonsmooth and $K = \phi$, resulted in the *fully robust OSCV* method proposed by Savchuk et al. (2013).

3 Fully Robust OSCV

In the fully robust OSCV method one uses the fact that the rescaling constants C and C^* are completely determined by the kernels K and H. For fixed K one may choose H such that $C = C^*$. A kernel H that produces such equality is called *robust* since it makes the OSCV method consistent regardless of smoothness of r. The E_{MASE} measure for a robust kernel is identical zero.

Savchuk et al. (2013) fixed $K = \phi$ and found a robust kernel in the following family:

$$H_I(x) = (1+\alpha)\phi(x) - \frac{\alpha}{\sigma}\phi\left(\frac{x}{\sigma}\right), \quad x, \alpha \in \mathbb{R},\ \sigma > 0. \tag{16}$$

The subscript "I" is used to indicate that the kernels (16) originate from the indirect cross-validation method of Savchuk et al. (2010). The robust kernel used in Savchuk et al. (2013) has

$$\begin{aligned} &\alpha = 0.0000879985198548436 \ \text{ and} \\ &\sigma = 10. \end{aligned} \tag{17}$$

The solution (17) was originally found in the way explained below.

The family (16) produces probability density functions for $-1 \le \alpha \le 0$ and $\sigma > 0$ or for $\alpha > 0$ and $\alpha/(1+\alpha) \le \sigma \le 1$. We did not find any nonnegative robust kernels in the range $-1 \le \alpha \le 100$. We, thus, started to search for robust kernels in the region $\alpha > 0$ and $\sigma > 1$ that corresponds to the kernels with negative tails. Observe that $\alpha = 0$ yields $H_I \equiv \phi$. Since the constants C_ϕ and C_ϕ^* are close, we looked for insignificant modification of

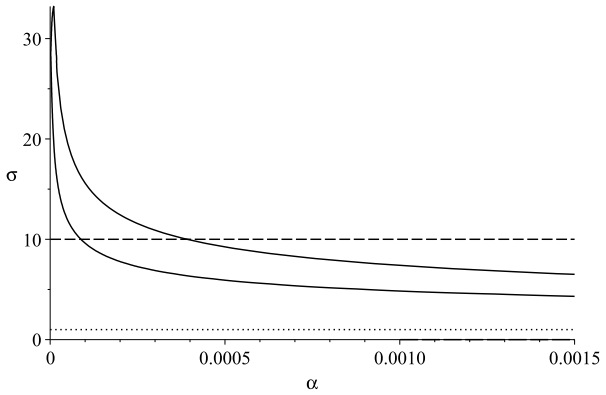

Figure 1: Robust negative-tailed kernels. The dotted and dashed lines correspond to $\sigma = 1$, and $\sigma = 10$, respectively.

ϕ that correspond to the values of α close to zero. Figure 1 shows the robust negative-tailed kernels in the range $0.00001 \leq \alpha \leq 0.0015$.

All kernels in Figure 1 are close to ϕ. Savchuk et al. (2013) arbitrary selected $\sigma = 10$ and used the solution (17). The other solution at $\sigma = 10$ has $\alpha = 0.0003912884532000514$. The corresponding kernel has somewhat larger L^2 distance compared to the kernel defined by (17), but still performs almost identical to it. Actually, all robust kernels shown in Figure 1 are really close and perform similarly. In what follows, we concentrate on using H_I with the values of the parameters as in (17).

The kernel H_I has the unique rescaling constant $C_I = 0.5217$ that is appropriate in both smooth and nonsmooth cases. Let $\hat{b}_I$ denote the minimizer of the OSCV function (10) that is computed based on H_I. The corresponding bandwidth that is used to compute a regression estimate is $\hat{h}_I = C_I \hat{b}_I$. In what follows, $\hat{b}_{OSCV}$ corresponds to the minimizer of the OSCV curve based on the Gaussian kernel ϕ, and $\hat{h}_{OSCV} = C_\phi \hat{b}_{OSCV}$.

The kernels ϕ and H_I look virtually the same when plotted on the interval $[-4, 4]$. It turns out that $H_I(x) < 0$ for $|x| > 4.85$. It is also remarkable that the tails of $H_I(x)$ are close to zero even for "large" x. Thus, $H_I(\pm 10) = -2.13 \cdot 10^{-6}$. It appears that the efficiency of H_I, computed according to Wand and Jones (1995), is 0.9552, that is even larger than 0.9512 in the case of ϕ. Nevertheless, we only use H_I for the cross-validation purposes since it has negative tails.

Let L_ϕ and L_I denote the one-sided counterparts of ϕ and H_I, respectively, computed according to (9). Closeness of L_ϕ and L_I on the interval $[0, 4]$ even seems to contradict the fact that $C_\phi = 0.6168$, whereas $C_I = 0.5217$. It follows from (11) that the discrepancy in C_ϕ and C_I is caused by the difference in the constants J_{L_ϕ} and J_{L_I}, obtained from L_ϕ and L_I according to (2). Indeed, $J_{L_\phi} = 1.2586$, whereas $J_{L_I} = 1.4882$. Consider the numerical values of the constituents of J_{L_ϕ} and J_{L_I}:

$$\begin{aligned} R_{L_\phi} &= 1.7860, \quad R_{L_I} = 1.8230, \\ \mu^2_{2L_\phi} &= 0.5654, \quad \mu^2_{2L_I} = 0.2497. \end{aligned}$$

The values R_{L_ϕ} and R_{L_I} are quite close. It appears that the squared second moment is the culprit in causing the discrepancy between J_{L_ϕ} and J_{L_I}. The mismatch in $\mu^2_{2L_\phi}$ and $\mu^2_{2L_I}$ must be explained by different behaviour of L_ϕ and L_I in the tails. Define the following

integrals for a one-sided kernel L:

$$M_L(t) = \int_0^t u^2 L(u)\, du,$$

$$F_L(t) = \left(\frac{\int_0^t L(u)^2\, du}{M_L^2(t)} \right)^{1/5}.$$

Observe that $M_L(t) \to \mu_{2L}$ and $F_L(t) \to J_L$ as $t \to \infty$. Figure 2 contains a plot of $M_{L_I}^2(t)$ for $4 \le t \le 50$. The dashed and the dotted horizontal lines indicate the values of $\mu_{2L_I}^2$ and $\mu_{2L_\phi}^2$, respectively. The curve $M_{L_\phi}^2(t)$ is not shown since it is indistinguishable from the dotted line for $t \ge 5$. The figure shows that $M_{L_I}(t)$ substantially deviates from $\mu_{2L_\phi}^2$ for t larger than about 15.

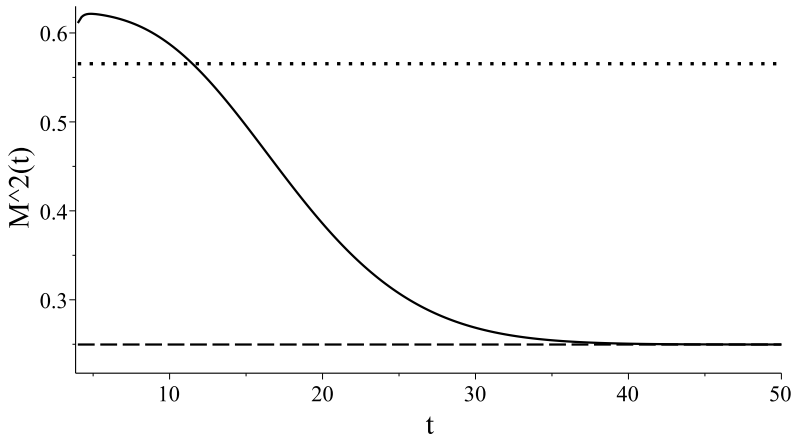

Figure 2: $M_{L_I}^2(t)$ for $4 \le t \le 50$. Dashed line shows $\mu_{2L_I}^2 = 0.2497$; dotted line shows $\mu_{2L_\phi}^2 = 0.5654$.

Figure 3 shows a plot of F_{L_I} for $4 \le t \le 50$. The graph of F_{L_ϕ} is not shown since for $t \ge 5$ it practically coincides with the dashed line showing J_{L_ϕ}. It turns out that $F_{L_I}(t) \ge 1.05 J_{L_\phi}$ for $t \ge 16.92$. This suggests that as long as L_I is not evaluated at a value larger than about 16.92, there is no reason to think that, in practical sense, using L_I is any different than using L_ϕ. It follows from (5), (6), and (10) that we evaluate L_I at values of the form $(x_i - x_j)/b$, where b is a bandwidth, and $x_1 < x_2 < \ldots < x_n$ are the design points on the interval $[0, a]$. Observe that $\max_{i,j} |x_i - x_j| \approx a$. In the case when r is smooth, the AMASE_{L_I}-optimal bandwidth, that is a proxy to the OSCV minimizer $\hat{b}_I$, is less than $a/16.92$ for

$$n > \left(\frac{16.92}{a} \right)^5 \frac{R_{L_I} \sigma^2}{\mu_{2L_I}^2 \int_0^a \left(r''(x) \right)^2 f(x)\, dx} = 1.0124 \cdot 10^7 \frac{\sigma^2}{a^5 \int_0^a \left(r''(x) \right)^2 f(x)\, dx}. \qquad (18)$$

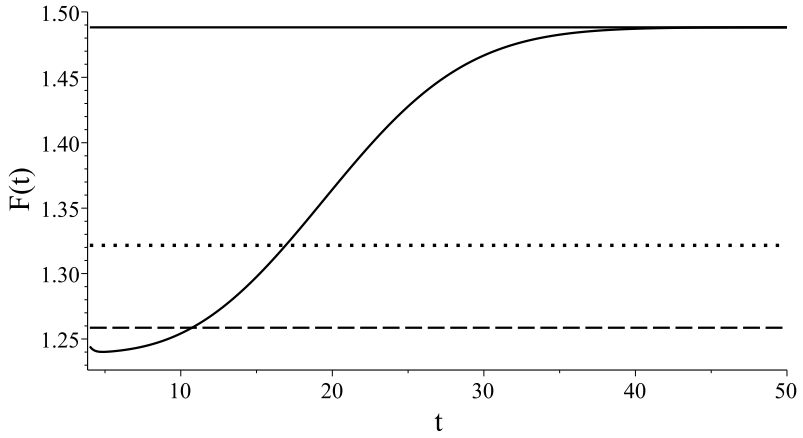

Figure 3: $F_{L_I}(t)$ for $4 \le t \le 50$. Dashed line shows $J_{L_\phi} = 1.2586$; solid horizontal line shows $J_{L_I} = 1.4882$; dotted line shows $1.05 J_{L_\phi}$.

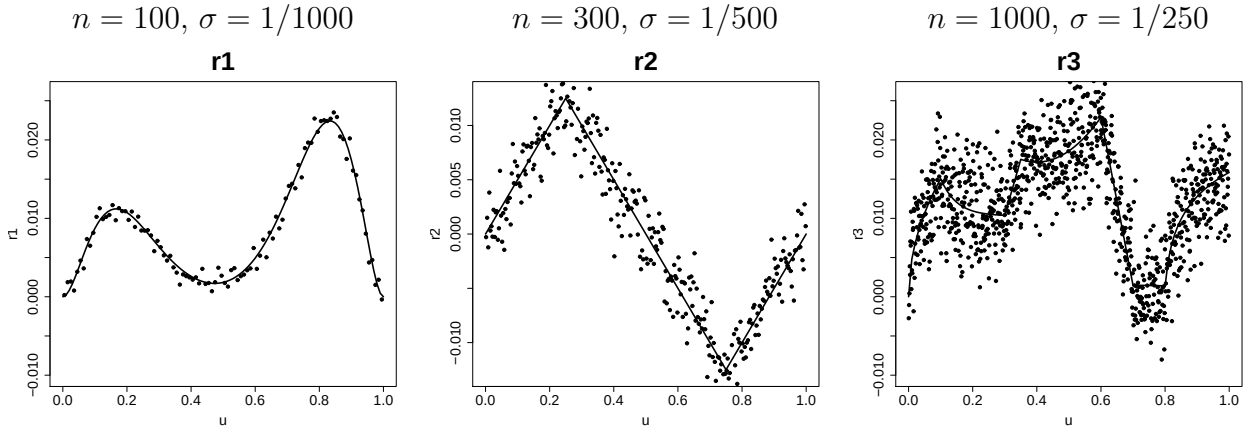

Figure 4: Regression functions r_1, r_2, r_3 and generated data.

In the nonsmooth case, the $\mathrm{AMASE}^*_{L_I}$-optimal bandwidth is less than $a/16.92$ given that

$$n > \left(\frac{16.92}{a}\right)^4 \frac{R_{L_I}\sigma^2}{3B_{L_I}\sum_{t=1}^{k} f(x^t)\left(r'(x^t+) - r'(x^t-)\right)^2} = 9.7408 \cdot 10^5 \frac{\sigma^2}{a^4 \sum_{t=1}^{k} f(x^t)\left(r'(x^t+) - r'(x^t-)\right)^2}. \quad (19)$$

In (18) and (19), the sample size is proportional to σ^2. This suggests that larger n is required for noisier data to make the difference between using L_ϕ and L_I essential.

For numerical illustration of (18) and (19), we use three regression functions, r_1, r_2, and r_3, that originate from the simulation study of Savchuk et al. (2013) and are defined in the Appendix. Figure 4 shows the graphs of r_1, r_2, and r_3 along with the typical data sets generated for specified n and σ.

Table 1 shows the smallest n that satisfies (18) for r_1 and (19) for r_2 and r_3 in the case $\sigma = 1/500$ and $f(x) = 1$. Since r_3 is the least smooth function of the three, it requires using smaller bandwidths and, consequently, involves the tail of L_I for smaller n compared to the

Table 1: Smallest n at which the difference between $\mu^2_{L_\phi}$ and $\mu^2_{2L_I}$ becomes essential in the case $f = 1$, and $\sigma = 1/500$.

Function	r_1	r_2	r_3
n	17	195	10

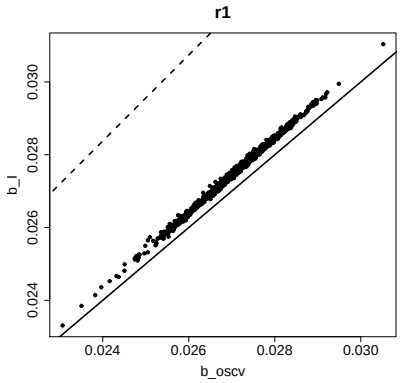

Figure 5: Scatter plots of $\hat{b}_I$ versus $\hat{b}_{OSCV}$ in the case of r_1, $n = 1000$, and $\sigma = 1/500$.

other two functions. The results in Table 1 indicate that the difference between using L_I and L_ϕ might be evident in finite samples.

The results of the numerical study of Savchuk et al. (2013) are used to compare the finite sample performances of the H_I-and ϕ-based OSCV versions. The sample sizes considered in Savchuk et al. (2013) are $n = 50$, 100, 300, and 1000, and the Gaussian noise levels are $\sigma = 1/250$, $1/500$, and $1/1000$. For a random variable Y defined in each replication of a simulation, let $\hat{E}(Y)$, $\hat{SD}(Y)$ and $\hat{M}(Y)$ denote the average, standard deviation, and median of Y over 1000 replications with r, n, and σ being fixed. One of the most important observations in the numerical study of Savchuk et al. (2013) is that $\hat{E}(\hat{h}_I/\hat{h}_{OSCV}) \approx \hat{SD}(\hat{h}_I/\hat{h}_{OSCV}) \approx C_I/C_\phi = 0.85$ for all considered regression functions, noise levels and sample sizes. This implies that $\hat{b}_I \approx \hat{b}_{OSCV}$. This result is not surprising in the nonsmooth case, where it is expected that for "large" n

$$\hat{b}_I \approx \frac{C^*_\phi}{C_I}\hat{b}_{OSCV} = 1.0128 \cdot \hat{b}_{OSCV}.$$

The corresponding large sample result in the smooth case is

$$\hat{b}_I \approx \frac{C_\phi}{C_I}\hat{b}_{OSCV} = 1.1823 \cdot \hat{b}_{OSCV}.$$

Thus, in the smooth case $\hat{b}_I$ is expected to be somewhat larger than $\hat{b}_{OSCV}$ for "large" n. Figure 5 shows the scatter plots of $\hat{b}_I$ versus $\hat{b}_{OSCV}$ in the case of r_1, $n = 1000$, and $\sigma = 1/500$. The solid line in the plot shows the 45 degrees line that passes through the origin. The dashed line passes through the origin and has the slope equal to 1.1823. The points on the graph form a line that lies between the solid and dashed lines, but substantially closer to the former one compared to the latter one. Larger sample size might be needed for the points to lie closer to the line with the slope of 1.1823.

The result $\hat{b}_I \approx \hat{b}_{OSCV}$ is a consequence of the fact that the H_I- and ϕ-based OSCV curves computed for the same data set are usually drastically close except in the neighborhood of zero, where the H_I-based curve might occasionally exhibit spurious bumps, as illustrated in the data examples in Section 4.

Since $C_\phi > C_I$, the result $\hat{b}_I \approx b_{OSCV}$ implies that the H_I-based OSCV version is expected to produce too low bandwidths in the smooth case. For assessment of the finite sample relative bandwidth bias by the H_I- and ϕ-based OSCV versions, we use the numerical data of Savchuk et al. (2013) to compute

$$\Delta_B = \frac{\hat{M}(\hat{h}) - \hat{M}(\hat{h}_0)}{\hat{M}(\hat{h}_0)} \cdot 100\%.$$

Table 2 contains the values of Δ_B in the cases $\hat{h} = \hat{h}_I$ and $\hat{h}_{OSCV}$ for all regression functions, $n = 100$, 300, 1000 and $\sigma = 1/500$. In the case of r_1, the ϕ-based OSCV method produces

Table 2: Values of Δ_B for r_1, r_2, and r_3 in the case $\sigma = 1/500$ and $n = 100$, 300, and 1000.

Method	H_I-based OSCV			ϕ-based OSCV		
Function	r_1	r_2	r_3	r_1	r_2	r_3
$n = 100$	-12.97	-8.33	-3.37	2.10	8.00	13.46
$n = 300$	-12.76	-9.45	-5.33	1.91	6.48	11.40
$n = 1000$	-12.77	-8.84	-5.65	1.48	6.87	10.87

bandwidths that are slightly biased upward, whereas the H_I-based OSCV version has the relative bandwidth bias of about -13%. In the case of r_3, the value of Δ_B for H_I-based OSCV is still negative but much closer to zero, whereas the ϕ-based OSCV method has $\Delta_B > 10\%$ for all considered sample sizes. The case of r_2 is intermediate: the values of Δ_B for ordinary OSCV and H_I-based OSCV are similar in magnitude but have opposite signs. Both versions of the OSCV method seem to be tricked by the function r_2 that has two cusps that can be easily masked by the data's noise, as it is illustrated in Figure 4.

The measure Δ_B can be thought of an empirical analog of E_C in the nonsmooth case. Observe that for the ϕ-based OSCV method, the values of Δ_B in Table 2 do not approach the theoretical result $E_C = 16.73\%$ even in the case of r_3 and $n = 1000$.

The wiggles in the H_I-based OSCV curve are shown to be caused by negativity of the tails of H_I. This inspired a new search of nonnegative robust kernel that resulted in several robust bimodal kernels. One of the kernels, H_B, is defined as

$$H_B(x) = 5\phi(10(x + \mu)) + 5\phi(10(x - \mu)), \tag{20}$$

where $\mu = 0.412071682$. The rescaling constant for H_B is $C_B = 0.1932$. We empirically found that bimodality of H_B is associated with producing OSCV curves with several local minima that are often of comparable sizes. This is illustrated in Figure 6 **(a)**, that shows a typical H_B-based OSCV curve in the case of r_1, $\sigma = 1/500$, $n = 100$ and the Uniform$(0, 1)$ design. The corresponding ϕ-based OSCV curve, shown in Figure 6 **(b)**, is smooth and has one local minimum.

The other newly found robust bimodal kernels perform similarly to H_B. This supports the empirically derived conclusion of Savchuk et al. (2013) that a "good" robust kernel should be unimodal and nonnegative. The problem of finding such a kernel is still open.

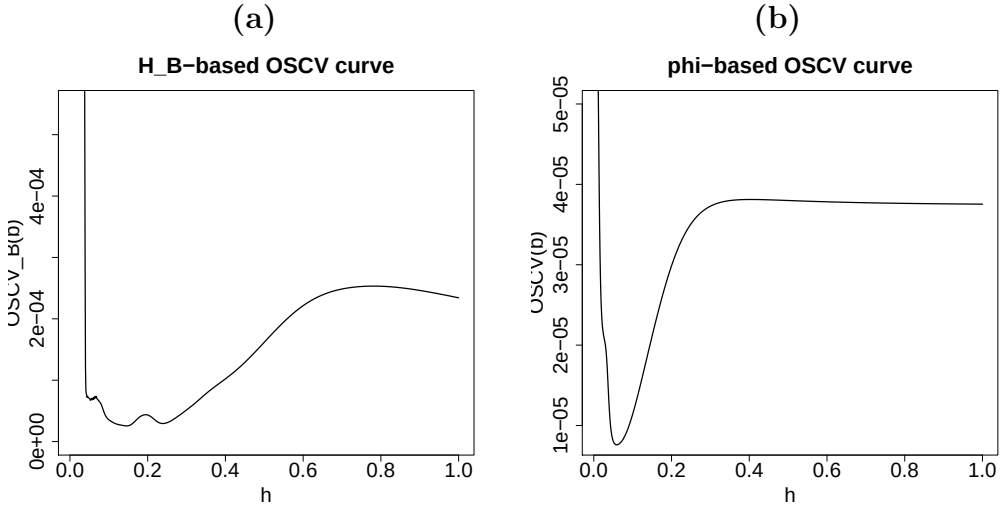

Figure 6: **(a)** H_B- and **(b)** ϕ-based OSCV curves for the data generated from r_1 in the case of $n = 100$, $\sigma = 1/500$, and the Uniform$(0, 1)$ design.

4 Data Examples

Performances of the H_I- and ϕ-based OSCV versions are further compared on the following two data examples.

4.1 Example 1 (Fuel consumption).

The data on car city-cycle fuel consumption in miles per gallon (mpg) can be downloaded from `http://archive.ics.uci.edu/ml/datasets/Auto+MPG` with the required citation of Lichman (2013). The same data set is used in Savchuk et al. (2013), but in this article we consider dependence of mpg (y) on car weight (x) instead of horsepower. Let $\delta_i = x_i - x_{i-1}$, $i = 2, \ldots, n$. Figure 7 shows the OSCV curves based on ϕ and H_I for $b > \min_i \delta_i$. For these data $\hat{b}_I \approx \hat{b}_{OSCV} \approx 490.74$. The OSCV curves based on H_I and ϕ are quite similar except for the values of b near zero, where the OSCV curve based on H_I has spurious wiggles, whereas the OSCV curve based on ϕ is smooth. Let $\hat{h}_{PI}$ denote the Ruppert-Sheather-Wand plug-in bandwidth computed for a given data set. The bandwidths selected by different methods for the data on fuel consumption are shown in the table below.

$\hat{h}_{OSCV}$	$\hat{h}_I$	$\hat{h}_{CV}$	$\hat{h}_{PI}$
302.71	256.02	270.64	263.67

The local linear regression estimate based on $\hat{h}_I$ is shown in Figure 8. The estimates based on the other bandwidths from the above table are similar.

The H_I-based OSCV curve in Figure 7 **(a)** behaves almost like a discontinuous function for "small" b. Alternating sign of $H_I((x_i - x_j)/b)$ for $x_j < x_i$, $i, j = 1, \ldots, n$, is one of the factors that occasionally produces very "small" sum of weights in the denominator of $\tilde{r}_b^i(x_i)$ for certain values of i and b. The resulting "large" value of $\tilde{r}_b^i(x_i)$ produces "large" squared deviation in the OSCV function (10) that causes a spike in the OSCV curve at the corresponding value of b.

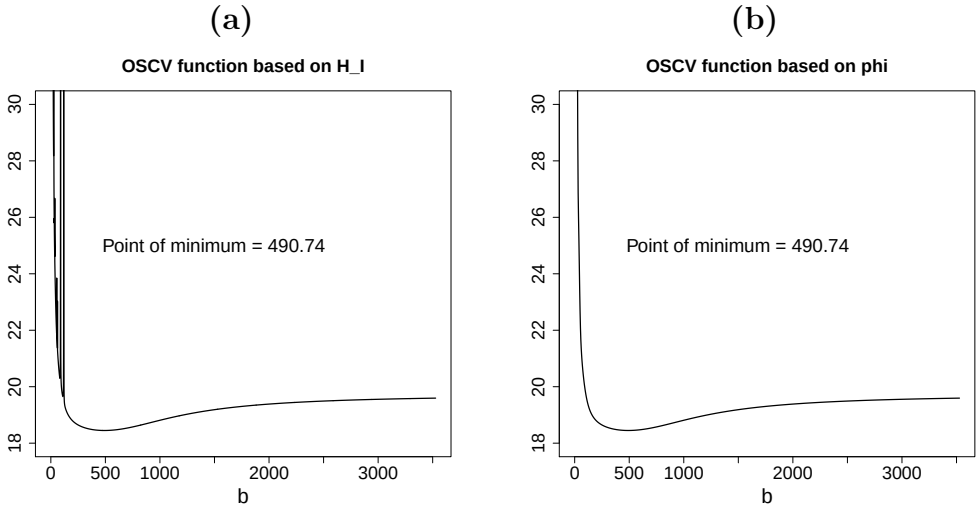

Figure 7: **(a)** H_I- and **(b)** ϕ-based OSCV curves for the mpg and car weight data.

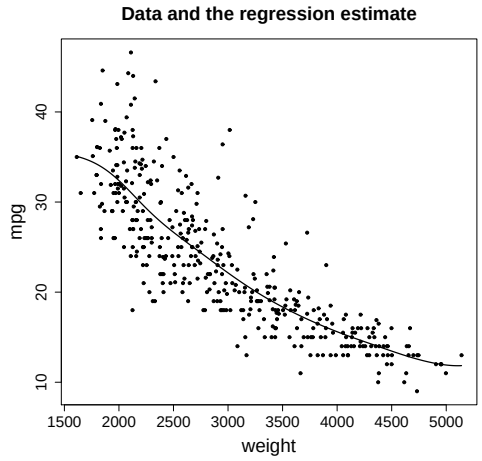

Figure 8: Regression estimate based on $\hat{h}_I = 256.02$ for the data on mpg and car weight.

The mpg and car weight example illustrates a typical behaviour of H_I-based OSCV for a data set of size $n \geq 100$. Spurious wiggles in the H_I-based OSCV curve usually appear for "small" b and do not interfere with the problem of determining $\hat{b}_I$. For $n < 100$ the wiggles may occasionally produce a fake global minimum of the H_I-based OSCV curve, as it is illustrated by the example in the following section.

4.2 Example 2 (Weight of rabbits).

Dudzinski and Mykytowycz (1961) studied the relationship between the eye lens weight and age of rabbits in Australia. The data set of size $n = 71$ can be downloaded from "http://www.statsci.org/data/oz/rabbit.html". Dudzinski and Mykytowycz (1961) constructed a model that relates the lens weight (y) to age (x) as

$$y = \alpha \exp\{-\beta/(x+\gamma)\}$$

for certain values of α, β, and γ. To the contrary of a parametric approach of Dudzinski and Mykytowycz (1961), we estimated r by using the LLE. The table below shows the bandwidths produced for the rabbits' data by different methods.

$\hat{h}_{OSCV}$	$\hat{h}_I$	$\hat{h}_{CV}$	$\hat{h}_{PI}$
50.34	23.42	46.95	54.48

All methods but fully robust OSCV produce comparable bandwidths and similar regression fits. Figure 9 **(a)** and **(b)** shows the H_I- and ϕ-based OSCV curves, correspondingly. For each graph the scale along the horizontal axis is changed such that the global minimum is attained at $\hat{h}_I = 23.42$ in the case of H_I and $\hat{h}_{OSCV} = 50.34$ in the case of ϕ. The cor-

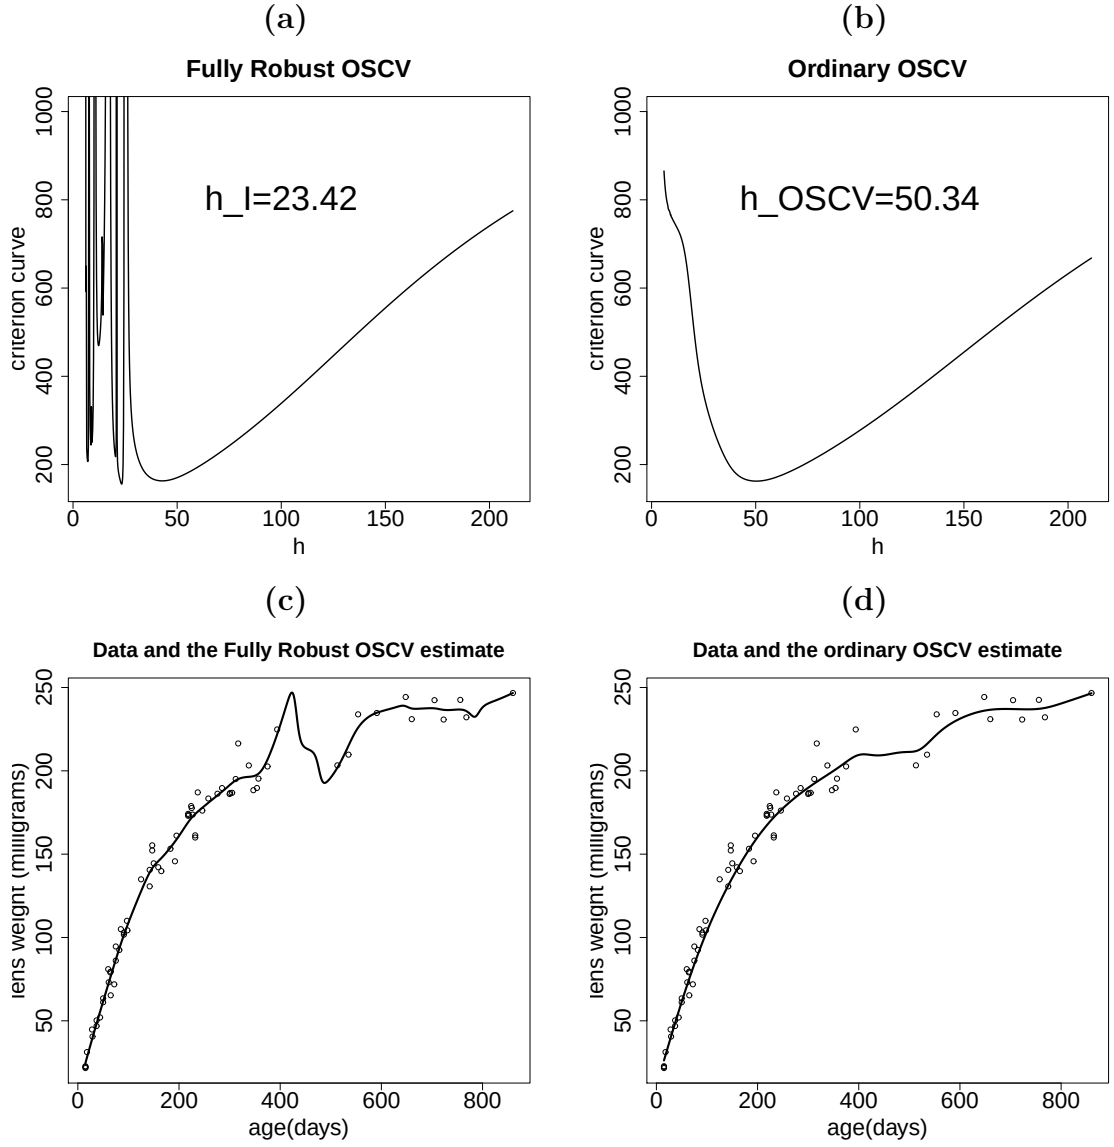

Figure 9: FROSCV and ordinary OSCV curves and regression fits for the data on eye lens weight and age.

responding local linear estimates are shown in Figure 9 **(c)** and **(d)**. The fit by ordinary

OSCV is quite similar to that obtained by Dudzinski and Mykytowycz (1961). The regression estimate produced by H_I is undersmoothed because of inappropriately small value of $\hat{h}_I$ obtained from a spurious wiggle of the H_I-based OSCV curve. Notice that the largest local minimum of the curve is attained at $h = 42.74$ that produces a regression estimate similar to the one corresponding to the ϕ-based OSCV version. This example and our numerous empirical experience suggest modifying the bandwidth section rule for the H_I-based OSCV version so that $\hat{h}_I$ corresponds to the largest local minimum of the H_I-based OSCV curve. This suggestion is similar to that given by Hall and Marron (1991) in the context of the kernel density estimation.

5 Summary and Conclusions

The OSCV method is a two-stage procedure. In the first stage one determines the minimizer of the OSCV curve computed based on the kernel H that is, generally, different from the kernel K used in computing the resulting regression estimate $\hat{r}_h$. The second stage consists in rescaling the bandwidth obtained in the first stage by using the multiplicative constant that is completely determined by K, H, and the smoothness of a regression function r. Unless the smoothness of r is specified, one by default uses the smooth rescaling constant, as this is the case in the original OSCV version of Hart and Yi (1998).

Out of the most often used kernels K, such as the Epanechnikov, quartic or Gaussian kernel, the latter one has the largest discrepancy between the smooth and nonsmooth rescaling constants. Thus, for $K = H = \phi$, using the smooth rescaling constant in the case of a nonsmooth function r results in the asymptotic relative bandwidth bias of 16.74%. Asymptotically, this bias further produces 4.02% MASE increase. This inspired Savchuk et al. (2013) to develop the method's correction, termed fully robust OSCV. The idea behind the fully robust OSCV method is to set $K = \phi$ and choose H that produces equal smooth and nonsmooth rescaling constants. Such a kernel H is called robust since it produces consistent OSCV bandwidths regardless of smoothness of r.

The current implementation of the fully robust OSCV method is based on the kernel H_I that is drastically close to the Gaussian kernel ϕ in a wide range of values of an argument, but has negative tails. Despite this fact, the second moments of L_ϕ and L_I, the one-sided counterparts of ϕ and H_I, respectively, are quite different. The discrepancy is caused by different tail behaviours of L_ϕ and L_I. This difference is the main factor that leads to equality of the smooth and nonsmooth rescaling constants in the case of H_I.

The practical performances of the H_I- and ϕ-based OSCV versions are compared based on the real data examples and the results of the numerical study of Savchuk et al. (2013). For a given data set, the H_I- and ϕ-based OSCV curves are usually quite close, except in the neighborhood of zero, where the H_I-based curve might exhibit spurious bumps, that are the artifacts of negativity of the tails of H_I. Except for small sample sizes ($n < 100$), where the wiggles in the H_I-based curve may result in a fake global minimum, the minimizers of the H_I-based and ϕ-based OSCV curves, $\hat{b}_I$ and $\hat{b}_{OSCV}$, respectively, are usually about the same in both smooth and nonsmooth cases. To avoid the problem of selecting $\hat{h}_I$ from the "wiggly part" of the H_I-based OSCV curve that might happen at "small" n, we suggest that $\hat{h}_I$ corresponds to the largest H_I-based OSCV curve minimizer.

In finite samples, the distribution of the H_I-based OSCV bandwidths is usually shifted downwards compared to that of the ASE-optimal bandwidths. However, the magnitude of

the relative bandwidth bias by H_I decreases as the smoothness of r decreases. We assessed the absolute value of the relative bandwidth bias produced by H_I and ϕ for $100 \leq n \leq 1000$. For H_I, the absolute value of the relative bandwidth bias is about 13% in the case of r_1, 9% in the case of r_2, and, finally, about 5% in the case of r_3. For ϕ, the relative bandwidth bias is under 2.1% in the case of r_1, but it exceeds 6% in the case of r_2 and 10% in the case of r_3. In the nonsmooth case, the asymptotically predicted relative bandwidth bias of 16.74% is not attained by ϕ in the considered range of n values, even in the case of the least smooth regression function r_3.

The relative bandwidth bias computation and the fact $\hat{b}_I \approx b_{OSCV}$ suggest that in the smooth case using H_I instead of ϕ is practically equivalent to adding wiggles to the OSCV curve along with using a wrong rescaling constant that produces too low bandwidth. There is some benefit of using H_I in the case where r has multiple cusps, though. However, since the nonsmooth Gaussian constant C_ϕ^* is about equal to C_I, we suggest that in the case when nonsmoothness of r is evident from the scatter diagram of the data, one uses ϕ along with C_ϕ^* instead of H_I along with C_I. The benefit of the former combination over the latter one is obtaining a smoother OSCV curve.

The kernel H_I uncovers the OSCV method's sensitivity to insignificant modifications of the kernel used in the cross-validation stage. Indeed, tiny deviation of H_I from ϕ in the tails greatly changes theoretical properties and practical performance of the OSCV method.

Nonnegative robust kernels are expected to produce smoother OSCV curves compared to the negative-tailed kernel H_I. Our search for nonnegative robust kernels resulted in the bimodal kernel H_B and several other robust bimodal kernels. Even though H_B yields smoother OSCV curves compared to H_I, we found that bimodality of H_B is associated with producing the curves with multiple local minima. This encourages a new search for *nonnegative unimodal* robust kernels in the case $K = \phi$.

References

W. S. Cleveland. Robust locally weighted regression and smoothing scatterplots. *J. Amer. Statist. Assoc.*, 74(368):829–836, 1979. ISSN 0003-1291.

M. Dudzinski and R. Mykytowycz. The eye lens as an indicator of age in the wild rabbit in australia. *CSIRO Wildlife Research*, 6:156–159, 1961.

T. Gasser and H.-G. Müller. Kernel estimation of regression functions. In *Smoothing techniques for curve estimation (Proc. Workshop, Heidelberg, 1979)*, volume 757 of *Lecture Notes in Math.*, pages 23–68. Springer, Berlin, 1979.

P. Hall and J. S. Marron. Local minima in cross-validation functions. *Journal of the Royal Statistical Society, Series B*, 53(1):245–252, 1991.

J. D. Hart and C.-L. Lee. Robustness of one-sided cross-validation to autocorrelation. *J. Multivariate Anal.*, 92(1):77–96, 2005. ISSN 0047-259X.

J. D. Hart and S. Yi. One-sided cross-validation. *Journal of the American Statistical Association*, 93(442):620–631, 1998.

M. Lichman. UCI machine learning repository, 2013. URL `http://archive.ics.uci.edu/ml`.

M. D. Martínez-Miranda, J. P. Nielsen, and S. Sperlich. One sided cross validation for density estimation. In G. N. Gregoriou, editor, *Operational Risk Towards Basel III: Best Practices and Issues in Modeling, Management and Regulation*, pages 177–196. John Wiley & Sons, Hoboken, New Jersey, 2009.

D. Ruppert, S. J. Sheather, and M. P. Wand. An effective bandwidth selector for local least squares regression. *J. Amer. Statist. Assoc.*, 90(432):1257–1270, 1995. ISSN 0162-1459.

O. Y. Savchuk, J. D. Hart, and S. J. Sheather. Indirect cross-validation for density estimation. *Journal of the American Statistical Association*, 105(489):415–423, 2010.

O. Y. Savchuk, J. D. Hart, and S. P. Sheather. One-sided cross-validation for nonsmooth regression functions. *J. Nonparametr. Stat.*, 25(4):889–904, 2013. ISSN 1048-5252.

C. J. Stone. Consistent nonparametric regression. *Ann. Statist.*, 5(4):595–645, 1977. ISSN 0090-5364. With discussion and a reply by the author.

M. P. Wand and M. C. Jones. *Kernel smoothing*, volume 60 of *Monographs on Statistics and Applied Probability*. Chapman and Hall Ltd., London, 1995. ISBN 0-412-55270-1.

S. Yi. A comparison of two bandwidth selectors OSCV and AICc in nonparametric regression. *Comm. Statist. Simulation Comput.*, 34(3):585–594, 2005. ISSN 0361-0918.

Appendix

Regression functions r_1, r_2, and r_3 are defined below. For each function, $0 \le x \le 1$.

$$r_1(x) = 5x^{10}(1-x)^2 + 2.5x^2(1-x)^{10},$$

$$r_2(x) = \begin{cases} 0.0125 - 0.05|x - 0.25|, & 0 \le x \le 0.5, \\ 0.05|x - 0.75| - 0.0125, & 0.5 < x \le 1. \end{cases}$$

$$r_3(x) = \begin{cases} 0.047619\sqrt{x}, & 0 \le x < 0.1, \\ 0.035186e^{-20x} + 0.010297, & 0.1 \le x < 0.3, \\ 0.142857x - 0.032473, & 0.3 \le x < 0.35, \\ 0.142857(x-0.35)(x-0.45) + 0.017527, & 0.35 \le x < 0.6, \\ 0.151455 - 0.214286x, & 0.6 \le x < 0.7, \\ 0.001455 - 0.214286(x-0.7)^3(x-0.4), & 0.7 \le x < 0.8, \\ 0.004762\ln(10x - 7.9) + 0.012334, & 0.8 \le x \le 1. \end{cases}$$